

Fast Color Space Transformations Using Minimax Approximations

M. Emre Celebi

Dept. of Computer Science
Louisiana State Univ., Shreveport, LA, USA
ecelebi@lsus.edu

Hassan A. Kingravi

Dept. of Computer Science, Georgia Institute of Technology, Atlanta, GA, USA
kingravi@gatech.edu

Fatih Celiker

Dept. of Mathematics, Wayne State University, Detroit, MI, USA
celiker@math.wayne.edu

September 7, 2010

Abstract

Color space transformations are frequently used in image processing, graphics, and visualization applications. In many cases, these transformations are complex nonlinear functions, which prohibits their use in time-critical applications. In this paper, we present a new approach called Minimax Approximations for Color-space Transformations (MACT). We demonstrate MACT on three commonly used color space transformations. Extensive experiments on a large and diverse image set and comparisons with well-known multidimensional lookup table interpolation methods show that MACT achieves an excellent balance among four criteria: ease of implementation, memory usage, accuracy, and computational speed.

Introduction

Color space transformations are commonly used in various image processing, graphics, and visualization tasks for decoupling luminance and chromaticity information, ensuring approximate perceptual uniformity, or achieving invariance to different imaging conditions such as viewing direction, illumination intensity, and highlights. Since color devices are usually provided with direct RGB signal input and output, the RGB color space is generally the source in these transformations.

In many cases, color space transformations are complex nonlinear functions, which prohibits their use in time-critical applications. In this paper, we present a new approach called Minimax Approximations for Color-space Transformations (MACT). We demonstrate MACT on transformations between the RGB space and three popular color spaces: CIELAB, HSI, and SCT. There are various reasons behind the choice of these particular spaces. First, all three spaces decouple luminance and chromaticity information, which makes them suitable for tasks including enhancement [1, 2], rendering [3], noise removal [4, 5], segmentation [6, 7, 8], and object recognition [9, 10]. Second, an important feature of CIELAB is its approximate perceptual uniformity, an essential requirement for expressing color differences in applications such as color quantization [11], mesh denoising [12], and maximum contrast color set design [13]. Third, HSI and SCT are designed to match human intuition, which makes them useful for manual color selection [14]. Last, the first two components of HSI are shown to be invariant to viewing direction, surface orientation, illumination direction, and illumination intensity [10].

The rest of the paper is organized as follows. Section 2 gives the related work. Section 3 presents the use of minimax approximation theory in speeding up color space transformations and the experimental results. Finally, Section 4 gives the conclusions.

2 Related Work

Traditionally, color space transformations in digital imaging systems have been implemented using lookup tables (LUTs) that require some form of multidimensional interpolation. The most commonly used 3D LUT interpolation methods are trilinear, prism, pyramidal, and tetrahedral [15]. These methods will be explained in §3.5.3. In this section, we briefly review the interpolation methods that appear less frequently in the literature as well as some alternative approaches.

Chang et al. [16] developed a method called Sequential Linear Interpolation (SLI) that uses a partially separable grid structure, which allows the allocation of more grid points to the regions where the function to be interpolated is more nonlinear.

This scheme often results in more accurate transformations at the expense of increased computational cost when compared to trilinear interpolation. Kacker et al. [17] proposed a wavelet based method that uses a multiscale grid structure. This method is shown to achieve lower maximum error but higher average error when compared to SLI. Gupta and Gray [18] presented a method called Maximum Entropy Estimation, which is a generalization of tetrahedral interpolation. This method is shown to be more accurate than tetrahedral interpolation in general while being at best only half as fast. Hemingway [19] described a method based on n -simplexes that is slightly faster than tetrahedral interpolation. However, the accuracy of the method is not discussed in the paper.

Neural networks have been applied to the color space transformation problem in a number of studies [20, 21, 22, 23]. They have the advantages of being more flexible and requiring less memory. However, they require training and parameter tuning, and are prone to overfitting [18].

3 Approximate Color Space Transformations

3.1 Overview of Minimax Approximation Theory

Given a function f , we would like to approximate it by another function g so that the error (ε) between them over a given interval is arbitrarily small. The existence of such approximations is stated by the following theorem:

Theorem 3.1 (Weierstrass) *Let f be a continuous real-valued function defined on $[a, b]$, i.e. $f \in C[a, b]$. Then $\forall \varepsilon > 0$ there exists a polynomial P such that $\|f - P\| < \varepsilon$, i.e. $\forall x \in [a, b]$, $|f(x) - P(x)| < \varepsilon$.*

This is commonly known as the minimax approximation to a function. It differs from other methods, e.g. least squares approximations, in that it minimizes the maximum error (ε) rather than the average error:

$$\varepsilon = \max_{x \in [a, b]} |f(x) - P(x)| \quad (1)$$

A similar theorem establishes the existence of a rational variant of this method [24]. Let $n \geq 0$ be a natural number and let

$$P_n([a, b]) = \{a_0 + a_1x + \dots + a_nx^n : x \in [a, b], a_i \in \mathbb{R}, i = 0, 1, \dots, n\} \quad (2)$$

be the set of all polynomials of degree less than or equal to n . The set of irreducible rational functions, $R_m^n([a, b])$, is defined as:

$$R_m^n([a, b]) = \left\{ \frac{p(x)}{q(x)} : p(x) \in P_n([a, b]), q(x) \in P_m([a, b]) \right\} \quad (3)$$

where p and q have no common factors. Then [24]:

Theorem 3.2 *For each function $f \in C[a, b]$, there exists at least one best rational approximation from the class $R_m^n([a, b])$.*

This theorem states the existence of a rational approximation $r^* \in R_m^n([a, b])$ to a function $f \in C[a, b]$ that is optimal in the Chebyshev sense:

$$\max_{x \in [a, b]} |f(x) - r^*(x)| = \text{dist}(f, R_m^n) \quad (4)$$

where $\text{dist}(f, R_m^n)$ denotes the distance between f and $R_m^n([a, b])$ with respect to some norm, in our case, the Chebyshev (maximum) norm. Regarding the choice between a polynomial and a rational approximant, it can be said that certain functions can be approximated more accurately by rationals than by polynomials. Jean-Michel Muller explains this phenomenon as follows “It seems quite difficult to predict if a given function will be much better approximated by rational functions than by polynomials. It makes sense to think that functions that have a behavior that is ‘highly nonpolynomial’ (finite limits at $\pm\infty$, poles, infinite derivatives, ...) will be poorly approximated by polynomials.” [25].

In this study, the Remez Exchange Algorithm is used to calculate the minimax approximations. The reader is referred to [24, 25] for more information on the theory of minimax approximations and [26] for the implementation details of the Remez algorithm.

3.2 CIELAB Color Space

3.2.1 Color Space Description

CIELAB is an approximately uniform color space standardized by the CIE (Commission Internationale de l’Eclairage) in 1976 [27]. The L^* component represents the lightness, whereas a^* and b^* represent the chromaticity. The transformation between the RGB and CIELAB color spaces is comprised of three steps: (1) conversion from nonlinear R’G’B’ to linear RGB, (2) conversion from linear RGB to CIEXYZ, and (3) conversion from CIEXYZ to CIELAB.

In order to convert the nonlinear R'G'B' values to the linear RGB ones, inverse gamma correction (ITU-R BT.709) is performed:

$$k' \in \left\{ \frac{R'}{255}, \frac{G'}{255}, \frac{B'}{255} \right\}$$

$$k \in \{R, G, B\} = \begin{cases} \frac{k'}{4.5} & 0 \leq k' < 0.081 \\ \left(\frac{k' + 0.099}{1.099} \right)^{\frac{1}{0.45}} & 0.081 \leq k' \leq 1 \end{cases} \quad (5)$$

The conversion from linear RGB to CIEXYZ (ITU-R BT.709) is given by:

$$\begin{aligned} X &= 0.412391R + 0.357584G + 0.180481B \\ Y &= 0.212639R + 0.715169G + 0.072192B \\ Z &= 0.019331R + 0.119195G + 0.950532B \end{aligned} \quad (6)$$

Finally, the conversion from CIEXYZ to CIELAB is given by:

$$\begin{aligned} L^* &= 116f(Y/Y_0) - 16 \\ a^* &= 500[f(X/X_0) - f(Y/Y_0)] \\ b^* &= 200[f(Y/Y_0) - f(Z/Z_0)] \\ f(t) &= \begin{cases} t^{1/3} & t > 0.008856 \\ 7.787t + 16/116 & t \leq 0.008856 \end{cases} \end{aligned} \quad (7)$$

Here, X_0 , Y_0 , and Z_0 are the tristimulus values of the reference white. For the illuminant D65 these are:

$$\begin{aligned} X_0 &= 0.950456 \\ Y_0 &= 1.0 \\ Z_0 &= 1.089058 \end{aligned} \quad (8)$$

The distance between two pixels \mathbf{x} and \mathbf{y} in the CIELAB space is given by:

$$D_{CIELAB}(\mathbf{x}, \mathbf{y}) = \sqrt{(L_{\mathbf{x}}^* - L_{\mathbf{y}}^*)^2 + (a_{\mathbf{x}}^* - a_{\mathbf{y}}^*)^2 + (b_{\mathbf{x}}^* - b_{\mathbf{y}}^*)^2} \quad (9)$$

3.2.2 Approximation of the Cube-Root Function

(5) and (6) can be implemented efficiently using LUTs as commonly seen in the literature [15]. The cube-root function (*cbrt*) in (7) is the main factor that influences the computational time of the transformation. The probabilities of calling this function can be calculated from an image that contains every possible color in the 24-bit RGB space (see §3.5.1):

$$\begin{aligned} P(X/X_0 > 0.008856) &= 0.999626 \\ P(Y/Y_0 > 0.008856) &= 0.999165 \\ P(Z/Z_0 > 0.008856) &= 0.997184 \end{aligned} \quad (10)$$

These high values suggest that we can accelerate the transformation substantially if we can devise a fast approximation for the cube-root function (see Figure 1). Table 1 shows the coefficients of the minimax polynomials of various degrees. Here, n and ε_{max}^A represent the degree of the polynomial and error of the minimax approximation, respectively.




Figure 1: Cube-root function in the interval $[0.008856, 1]$

Table 1: Minimax polynomials for the cube-root function

n	ε_{max}^A	a_0	a_1	a_2	a_3	a_4	a_5
2	1.271154e-01	1.268979e-01	2.393873	-1.647669			
3	9.787829e-02	9.787826e-02	4.057495	-7.388864	4.331370		
4	8.111150e-02	8.111133e-02	5.926004	-2.017165e+01	2.833070e+01	-1.324728e+01	
5	7.002956e-02	7.002910e-02	7.961214	-4.329352e+01	1.063182e+02	-1.135685e+02	4.358268e+01

Table 2: Minimax rationals for the cube-root function

n	m	ε_{max}^A	a_0	a_1	a_2	a_3	a_4
2	2	2.060996e-03	6.309655e-03 4.482646e-02	5.785782e-01 1.175862	1.591005 9.596879e-01		
2	3	7.210231e-04	2.500705e-03 1.978701e-02	3.447113e-01 8.542797e-01	1.942708 1.664540	-2.503267e-01	
3	2	5.931593e-04	1.776519e-03 1.432256e-02	2.632323e-01 6.779998e-01	1.751297 1.706260	3.836709e-01	
2	4	3.107735e-04	1.317899e-03 1.124254e-02	2.390113e-01 6.726679e-01	2.099395 2.184656	-7.511150e-01	2.229717e-01
3	3	1.858694e-04	4.370889e-04 3.912364e-03	9.526952e-02 2.954084e-01	1.252009 1.717143	1.302733 6.343408e-01	
4	2	2.334688e-04	7.589302e-04 6.644723e-03	1.519784e-01 4.506424e-01	1.663584 2.030622	8.368075e-01	-1.657269e-01
3	4	8.539863e-05	1.683667e-04 1.610864e-03	4.667675e-02 1.617974e-01	9.106812e-01 1.494070	1.810577 1.218468	-1.079451e-01
4	3	8.052920e-05	1.349673e-04 1.299250e-03	3.832079e-02 1.345420e-01	7.870174e-01 1.330358	1.799062 1.365369	2.071170e-01
4	4	3.856930e-05	3.927283e-05 4.022100e-04	1.392318e-02 5.414536e-02	4.114739e-01 8.221526e-01	1.734853 1.800167	8.679223e-01 3.513617e-01

It can be seen that the error values are quite high and as the approximation degree is increased, the accuracy doesn't improve significantly. This suggests that rational functions might be better suited for this approximation task. Table 2 shows the coefficients of the minimax rationals of various degrees. Here, each pair of adjacent rows corresponds to a rational function of a particular degree (n, m) in which the first and second rows represent the numerator and denominator, respectively. It can be seen that minimax rationals can accurately represent the cube-root function.

3.3 HSI Color Space

3.3.1 Color Space Description

HSI (Hue-Saturation-Intensity) is an intuitive alternative to the RGB space [27]. It uses approximately cylindrical coordinates, and is a non-linear deformation of the RGB color cube. The hue H is a function of the angle in the polar coordinate system and describes a pure color. The saturation S is proportional to radial distance and denotes the purity of a color. Finally, the intensity I is the distance along the axis perpendicular to the polar coordinate plane and represents the brightness.

The transformation between RGB and HSI is given by:

$$\begin{aligned}
 H &= \arccos \left[\frac{0.5(R-G+R-B)}{\sqrt{(R-G)^2 + (R-B)(G-B)}} \right] \\
 \text{if } (B > G) \\
 H &= 2\pi - H \\
 S &= 1 - 3 \min(R, G, B) / (R + G + B) \\
 I &= (R + G + B) / 3
 \end{aligned} \tag{11}$$

where \arccos denotes the inverse cosine function. Kender [28] proposed a fast version of (11) that gives numerically identical results. This transformation involves fewer multiplications and no square root operation:

$$\begin{aligned}
 &\text{if } R > B \text{ and } G > B \\
 &\quad H = \frac{\pi}{3} + \arctan \left[\frac{\sqrt{3}(G-R)}{G-B+R-B} \right] \\
 &\text{else if } G > R \\
 &\quad H = \pi + \arctan \left[\frac{\sqrt{3}(B-G)}{B-R+G-R} \right] \\
 &\text{else if } B > G \\
 &\quad H = \frac{5\pi}{3} + \arctan \left[\frac{\sqrt{3}(R-B)}{R-G+B-G} \right] \\
 &\text{else if } R > B \\
 &\quad H = 0 \\
 &\text{else} \\
 &\quad H = \text{undefined}
 \end{aligned} \tag{12}$$

Table 3: Minimax polynomials for the inverse tangent function

n	ε_{max}^A	a_0	a_1	a_2	a_3	a_4	a_5
2	6.907910e-03	5.959793e-03	1.782975	7.497879e-01			
3	3.654156e-03	-3.654076e-03	1.884080	-9.805583e-01	1.430580e-01		
4	1.286371e-03	-1.286369e-03	1.796716	-4.958969e-01	-6.927404e-01	4.421541e-01	
5	1.801311e-04	-1.801283e-04	1.739333	-2.039848e-02	-2.065512	2.052837	-6.591729e-01

where \arctan denotes the inverse tangent function. The distance between two pixels \mathbf{x} and \mathbf{y} in the HSI space is given by:

$$D_{HSI}(\mathbf{x}, \mathbf{y}) = \sqrt{s_{\mathbf{x}}^2 + s_{\mathbf{y}}^2 - 2s_{\mathbf{x}}s_{\mathbf{y}}\cos\theta + (i_{\mathbf{x}} - i_{\mathbf{y}})^2}$$

$$\theta = \begin{cases} |h_{\mathbf{x}} - h_{\mathbf{y}}| & \text{if } |h_{\mathbf{x}} - h_{\mathbf{y}}| \leq \pi \\ 2\pi - |h_{\mathbf{x}} - h_{\mathbf{y}}| & \text{otherwise} \end{cases} \quad (13)$$


3.3.2 Approximation of the Inverse Tangent Function

We decided to approximate Kender's transformation (12) rather than the original one (11) due to several reasons. First, these two give identical results. Second, (12) is computationally cheaper than (11). Third, as will be seen in the next subsection, the inverse tangent function in (12) is easier to approximate when compared to the inverse cosine function in (11).

In (12), all of the cases involve multiplication by the constant $\sqrt{3}$. Therefore, a multiplication operation can be avoided by approximating $\arctan(\sqrt{3}x)$. Note that the argument of this function can also be negative. This can be handled using the following identity:

$$\arctan(x) = -\arctan(-x) \quad (14)$$

In (12), the inverse tangent function receives its arguments from the interval $[-1.0, 1.0]^1$. Figure 2 shows a plot of the function in the second half of this interval.

Figure 2: Inverse tangent function in the interval $[0, 1]$

Owing to its highly linear behavior, this function can be accurately approximated by low-order polynomials. Table 3 shows the coefficients of the minimax polynomials of various degrees (n).

3.4 Spherical Coordinate Transform (SCT) Color Space

3.4.1 Color Space Description

The Spherical Coordinate Transform (SCT) is defined as [29]:

$$\begin{aligned} L &= \sqrt{R^2 + G^2 + B^2} \\ \angle A &= \arccos(B/L) \\ \angle B &= \arccos\left(\frac{R}{L \sin(\angle A)}\right) \end{aligned} \quad (15)$$

where L represents the luminance, and angles $\angle A$ and $\angle B$ represent the chromaticity.

The expression for $\angle B$ can be simplified using trigonometric manipulations:

$$\angle B = \arctan(G/R) \quad (16)$$

¹This holds prior to the multiplication with $\sqrt{3}$.

This formulation is computationally advantageous in that it avoids a multiplication and a sine operation.

Unfortunately, it is not easy to define a perceptual distance function in SCT. Although various formulae have been developed to calculate the distance between two points lying on the same spherical surface, these cannot be used in SCT. This is because pixels of different brightness in this space lie on different spherical shells. Therefore, we decided to use the following alternative approach. Given two pixels in SCT, the inverse transformation (17) is applied to switch back to the RGB space. The pixels are then converted to CIELAB so that the distance between them can be calculated using (9). This indirect method of distance calculation is likely to introduce additional errors. However, minimax approximations for the elementary functions involved in (15) and (16) can be devised to obtain an arbitrarily accurate approximate transformation.

$$\begin{aligned} R &= L \sin(\angle A) \cos(\angle B) \\ G &= L \sin(\angle A) \sin(\angle B) \\ B &= L \cos(\angle A) \end{aligned} \quad (17)$$

3.4.2 Approximation of the Inverse Cosine Function

In the expression for $\angle A$, the inverse cosine function receives its arguments from the interval $[0, 1]$. Figure 3 shows a plot of this function.




Figure 3: Inverse cosine function in the interval $[0, 1]$

Unfortunately, approximating the inverse cosine function in this interval is not easy because of its behavior near 1. This can be circumvented using the following numerically more stable identity for $x \geq 0.5$:

$$\arccos(x) = 2 \arcsin\left(\sqrt{0.5(1-x)}\right) \quad (18)$$




Figure 4: Inverse sine function in the interval $[0, 0.5]$

In (18), the inverse sine function (\arcsin) receives its arguments from the interval $[0, 0.5]$. Figure 4 shows a plot of this function. In order to avoid two multiplication operations, the following function can be approximated instead:

$$\begin{aligned} y &= \sqrt{1-x} \\ \arccos(x) &= 2 \arcsin(y/\sqrt{2}) \end{aligned} \quad (19)$$

Table 4: Minimax polynomials for the inverse sine function

n	ε_{max}^A	a_0	a_1	a_2	a_3	a_4	a_5
4	2.097814e-05	2.097797e-05	1.412840	1.429881e-02	6.704361e-02	6.909677e-02	
5	2.370540e-06	-2.370048e-06	1.414434	-3.300037e-03	1.354670e-01	-3.994259e-02	6.099502e-02

Table 5: Minimax polynomials for the inverse cosine function


m	ε_{max}^A	a_0	a_1	a_2	a_3	a_4	a_5
4	1.048949e-05	1.570786	-9.990285e-01	-1.429899e-02	-9.481335e-02	-1.381942e-01	
5	1.186403e-06	1.570798	-1.000156	3.299810e-03	-1.915780e-01	7.988231e-02	-1.725177e-01

In (19), the argument y falls into the interval $[0, 1/\sqrt{2}]$. Table 4 shows the coefficients of the minimax polynomials of various degrees (n).

On the other hand, it can be seen from Figure 3 that the inverse cosine function is highly linear in the interval $[0, 0.5]$ and can be accurately approximated by polynomials. Table 5 shows the coefficients of the minimax polynomials of various degrees (m).

3.4.3 Approximation of the Inverse Tangent Function

In (16), the inverse tangent function receives its arguments from the interval $[0, 255]$. Figure 5 shows a plot of this function. Note that the degenerate cases, i.e. $R = G = 0$ ($\angle B = \text{undefined}$) and $R = 0, G > 0$ ($\angle B = \pi/2$), are not reflected in the plot.

Figure 5: Inverse tangent function in the interval $[0, 255]$

It can be seen that the function exhibits high variability in this large interval. In order to obtain an accurate low-order polynomial approximation, the domain can be divided into two as follows:

$$x = G/R$$

$$\arctan(x) = \begin{cases} f(x) = \arctan(x) & \text{if } G < R \\ g(x) = \pi/2 - \arctan(1/x) & \text{otherwise} \end{cases} \quad (20)$$

The second part of (20) follows from:

$$\arctan(x) = \pi/2 - \arctan(1/x) \quad \text{for } x > 0 \quad (21)$$

When $G < R$, the function f receives its arguments from the interval $[0, 254/255]$. On the other hand, when $R \leq G$ the inverse tangent function within g receives its arguments from the interval $[1/255, 1]$. Tables 6 and 7 show the coefficients of the minimax polynomials of various degrees (r) for the functions f and g , respectively.

As in the case of (10), the probabilities of calling the inverse sine and inverse cosine functions can be calculated from an image that contains every possible color in the 24-bit RGB space:

$$P_{\arcsin} = P(B/L \geq 0.5) = 0.557723$$

$$P_{\arccos} = 1 - P_{\arcsin} = 0.442277 \quad (22)$$

In contrast, the inverse tangent function in (16) is almost always called (except for when $R = G = 0$, in which case $\angle B = \text{undefined}$):

$$P_{\arctan} = 1 - P(R = G = 0) = 1 - \frac{256}{256^3} = 0.999985 \quad (23)$$

Note that $\angle B = \pi/2$ when $R = 0$ and $G > 0$.

Table 6: Minimax polynomials for the function f

r	ε_{max}^A	a_0	a_1	a_2	a_3	a_4	a_5
4	1.036515e-04	-1.036508e-04	1.003740	-1.773538e-02	-3.390563e-01	1.386796e-01	
5	2.073939e-05	2.073866e-05	9.982666e-01	2.352573e-02	-4.506862e-01	2.635050e-01	-4.920822e-02

Table 7: Minimax polynomials for the function g

r	ε_{max}^A	a_0	a_1	a_2	a_3	a_4	a_5
4	1.051643e-04	1.570917	-1.004004	1.885694e-02	3.373159e-01	-1.377930e-01	
5	2.012104e-05	1.570769	-9.981253e-01	-2.440212e-02	4.528921e-01	-2.659181e-01	5.016228e-02

3.5 Experimental Results

3.5.1 Calculation of the Color Space Transformation Errors

In order to calculate the accuracy of the presented approximate color space transformations, we used a $4,096 \times 4,096$ image (henceforth referred to as *RGB16Million*)², which contains 16,777,216 unique colors, i.e. every possible color in the 24-bit RGB space. Tables 8-10 show the average execution time (in seconds)³ over 1,000 identical runs and the average (ε_{avg}^T) and maximum (ε_{max}^T) transformation errors for CIELAB (9), HSI (13), and SCT (9) & (17), respectively. Note that the rows of Table 8 are sorted on $(n + m)$, since the execution time is proportional to the total degree of the polynomials in the rational. On the other hand, the rows of Table 10 are sorted on the degree of the inverse tangent approximation (r). This is because the probability of calling this function (23) is much higher than that of the other two elementary functions (22). Therefore, the execution time is mainly influenced by the degree of the inverse tangent approximation.

We also performed the exact transformations on *RGB16Million* and calculated the average execution time over 100 runs. The results were 60.331s, 13.485s, and 27.599s for CIELAB, HSI, and SCT, respectively. Comparing these values with those given in Tables 8-10, we can see that the proposed approximations provide substantial computational savings.

3.5.2 Calculation of the Computational Gain Values

It can be seen from Tables 8-10 that only marginal computational gains can be obtained using lower-order approximations without significantly compromising the accuracy of the transformation. Therefore, in this subsection, we consider only the highest-order approximations for each transformation. However, lower-order approximations might be preferable depending on the application requirements.

In order to calculate the computational gain for each approximate transformation, a set of 100 high quality RGB images was collected from the Internet. The set includes images of people, animals, plants, buildings, aerial maps, man-made objects, natural scenery, paintings, sketches, as well as scientific, biomedical, synthetic images and test images commonly used in the literature.

On each image, the exact transformations were performed 100 times and the average execution times were calculated. The same was done for the corresponding approximate transformations with 1,000 runs. The computational gain is calculated as the ratio of the average execution times. Table 11 shows the statistics over the entire image set. It can be seen that the computational gain observed in *RGB16Million* also applies to the case of real-world images.

3.5.3 Comparison with 3D Lookup Table Interpolation Methods

In this subsection we compare MACT with the most commonly used LUT interpolation methods, namely trilinear, prism, pyramidal, and tetrahedral. These methods involve three steps: packing, extraction, and interpolation [15]. Packing is a

²Available at <http://brucelindbloom.com/downloads/RGB16Million.tif.zip>

³Programming language: C, Compiler: gcc 3.4.4, CPU: Intel Pentium D 2.66Ghz

Table 8: Comparison of CIELAB approximations

n	m	Avg. Time	ε_{avg}^T	ε_{max}^T
2	2	1.031872	0.787741	2.225399
2	3	1.134568	0.295577	0.774404
3	2	1.114776	0.239033	0.636929
2	4	1.221390	0.131411	0.331015
3	3	1.165900	0.061127	0.193565
4	2	1.197828	0.090846	0.247291
3	4	1.291526	0.018092	0.085046
4	3	1.291570	0.014927	0.079935
4	4	1.292872	0.003201	0.036481

Table 9: Comparison of HSI approximations

n	Avg. Time	ε_{avg}^T	ε_{max}^T
2	0.865328	0.002523	0.007134
3	0.865575	0.001259	0.003693
4	0.866254	0.000448	0.001320
5	0.873251	0.000063	0.000190

Table 10: Comparison of SCT approximations

n	m	r	Avg. Time	ε_{avg}^T	ε_{max}^T
4	4	4	1.330927	0.006596	0.024273
4	5	4	1.371811	0.006553	0.024273
5	4	4	1.377924	0.006266	0.022820
5	5	4	1.419315	0.006224	0.021810
4	4	5	1.427337	0.002139	0.008371
4	5	5	1.462627	0.002011	0.008371
5	4	5	1.471486	0.001381	0.005793
5	5	5	1.507812	0.001254	0.004543

process that divides the domain of the source space and populates it with sample points to build the LUT. The extraction step aims at finding the location of the input pixel and extracting the color values of the nearest lattice points. The last step is 3D interpolation, in which the input point and the extracted lattice points are used to calculate the destination color specifications. Essentially, 3D interpolation is a repeated application of linear interpolation [15]:

$$x = (1 - \alpha)x_0 + \alpha x_1 \quad (24)$$

where x_0 and x_1 are the spatial coordinates of the two known points and α is the interpolation coefficient. The 3D interpolation is the step in which the aforementioned LUT interpolation methods differ. Trilinear interpolation uses 8 neighboring lattices, whereas prism, pyramidal, and tetrahedral interpolations use 6, 5, and 4 neighbors, respectively. The greater the number of neighboring lattices used in a method, the higher the computational requirements and accuracy. Since the formulations of these methods are mathematically involved, the interested reader is referred to the relevant literature [15].

We implemented two versions of each LUT method: standard and caching. The former is a direct implementation of the mathematical formulation, whereas the latter pre-computes the differences between neighboring lattices and stores these in each node. Table 12 shows the average execution times on *RGB16Million* over 1,000 runs and the transformation errors for CIELAB. It can be seen that even though these methods range from being 1.32 to 2.87 times faster than MACT, their accuracy is 18 to 161 times lower. In addition, these methods require extra storage of up to 10.7 MBs. Note that LUT size only affects the accuracy of the transformation and the storage requirements and, in theory, it should not affect the computational time. However, it can be observed from the table that as the LUT size is increased, the computational time increases as well. This is most likely due to the limited size of the processor cache.

As for HSI and SCT, since they include angular components (H component in HSI, $\angle A$ and $\angle B$ components in SCT), 3D interpolation in these color spaces involves the interpolation of angular (circular) data given by [30]:

$$\theta = \arctan \left(\frac{(1 - \alpha) \sin(\theta_0) + \alpha \sin(\theta_1)}{(1 - \alpha) \cos(\theta_0) + \alpha \cos(\theta_1)} \right) \quad (25)$$

where θ_0 and θ_1 are the two known angles and α is the interpolation coefficient. Due to the trigonometric and inverse-trigonometric functions in (25), it turns out that 3D interpolation in these spaces is computationally more expensive than the original transformations given in (11) and (15). For example, trilinear interpolation in HSI takes about 84s on *RGB16Million* ($\varepsilon_{avg}^T = 0.000777$, $\varepsilon_{max}^T = 1.086255$).

Note that even if the angular data was quantized into, say 360 steps, precomputing (25) would not be possible since the value of the interpolation coefficient α is not known *a priori*. In addition, the inverse tangent function cannot be approximated

Table 11: Computational gain statistics for each transformation

Transformation	Computational Gain				
	Min	Max	Mean	Stdev	Median
CIELAB	8.761736	40.629135	35.575373	4.608590	37.192729
HSI	4.511637	16.460248	11.437106	2.439482	11.669715
SCT	4.493543	22.232897	15.782453	2.110820	16.162434

Table 12: Comparison of 3D LUT interpolation methods

Method	LUT Size	Avg. Time (standard)	Avg. Time (caching)	ε_{avg}^T	ε_{max}^T
Trilinear	9^3	0.901673	0.619681	0.359677	5.409771
	17^3	0.908753	0.630465	0.100026	1.987040
	33^3	0.982880	0.716958	0.025089	0.653506
Prism	9^3	0.856652	0.609452	0.349128	5.409771
	17^3	0.860880	0.614681	0.095150	2.343905
	33^3	0.931349	0.677319	0.024172	0.826651
Pyramidal	9^3	0.850409	0.540481	0.319163	5.872591
	17^3	0.872813	0.566271	0.088116	2.909481
	33^3	0.940143	0.637049	0.022466	1.042665
Tetrahedral	9^3	0.689341	0.450413	0.284339	5.783921
	17^3	0.704121	0.486654	0.077870	2.702448
	33^3	0.773102	0.561076	0.020253	1.185788

using a minimax polynomial since its argument is not bounded. In short, the expensive interpolation operation offsets the computational advantage of 3D LUT interpolation in angular color spaces.

4 Conclusions

In this paper we proposed MACT, a novel approach to speed up color space transformations based on minimax approximations. Advantages of MACT include ease of implementation, negligible memory requirements, extremely good accuracy, and very high computational speed. Comparisons with commonly used 3D LUT interpolation methods revealed that MACT yields significantly more accurate transformations at the expense of slightly higher computational requirements. Although MACT was applied to three particular color space transformations, it can easily be adapted to other transformations, e.g. RGB to CIELUV transformation, that involve computationally expensive mathematical functions.

Implementations of the fast color space transformations described in this paper will be made publicly available at <http://sourceforge.net/projects/fourier-ipal>.

Acknowledgments

This publication was made possible by a grant from The Louisiana Board of Regents (LEQSF2008-11-RD-A-12). The authors are grateful to the anonymous reviewers for their valuable comments and to Bruce Lindbloom for providing the *RGB16Million* image.

References

- [1] W.-C. Kao and Y.-J. Chen, “Multistage Bilateral Noise Filtering and Edge Detection for Color Image Enhancement,” *IEEE Transactions on Consumer Electronics*, vol. 51, no. 4, pp. 1346–1351, 2005.
- [2] I. Pitas and P. Kiniklis, “Multichannel Techniques in Color Image Enhancement and Modeling,” *IEEE Transactions on Image Processing*, vol. 5, no. 1, pp. 168–171, 1996.
- [3] D.S. Ebert, C.J. Morris, P. Rheingans, and T.S. Yoo, “Designing Effective Transfer Functions for Volume Rendering from Photographic Volumes,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 8, no. 2, pp. 183–197, 2002.
- [4] L. Jin and D. Li, “A Switching Vector Median Filter Based on the CIELAB Color Space for Color Image Restoration,” *Signal Processing*, vol. 87, no. 6, pp. 1345–1354, 2007.
- [5] M.E. Celebi, H.A. Kingravi, B. Uddin, and Y.A. Aslandogan, “Fast Switching Filter for Impulsive Noise Removal from Color Images,” *Journal of Imaging Science and Technology*, vol. 51, no. 2, pp. 155–165, 2007.
- [6] D. Comaniciu and P. Meer, “Mean Shift: A Robust Approach toward Feature Space Analysis,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 5, pp. 603–619, 2002.
- [7] F. Perez and C. Koch, “Toward Color Image Segmentation in Analog VLSI: Algorithm and Hardware,” *International Journal of Computer Vision*, vol. 12, no. 1, pp. 17–42, 1994.

- [8] B.W. Minten, R.R. Murphy, J. Hyams, and M. Micire, “Low-Order-Complexity Vision-Based Docking,” *IEEE Transactions on Robotics and Automation*, vol. 17, no. 6, pp. 922–930, 2001.
- [9] M. Corbalán, M.S. Millán, and Yzuel M.J., “Color Pattern Recognition with CIELAB Coordinates,” *Optical Engineering*, vol. 41, no. 1, pp. 130–138, 2001.
- [10] T. Gevers and A.W.M. Smeulders, “Color Based Object Recognition,” *Pattern Recognition*, vol. 32, no. 3, pp. 453–464, 1999.
- [11] X. Wu, “Color Quantization by Dynamic Programming and Principal Analysis,” *ACM Transactions on Graphics*, vol. 11, no. 4, pp. 348–372, 1992.
- [12] S. Fleishman, I. Drori, and D. Cohen-Or, “Bilateral Mesh Denoising,” *ACM Transactions on Graphics*, vol. 22, no. 3, pp. 950–953, 2003.
- [13] C. Glasbey, G. van der Heijden, V. Toh, and A. Gray, “Colour Displays for Categorical Images,” *Color Research and Application*, vol. 32, no. 4, pp. 304–309, 2007.
- [14] S.-W. Hsiao, F.-Y. Chiu, and H.-Y. Hsu, “A Computer-Assisted Colour Selection System Based on Aesthetic Measure for Colour Harmony and Fuzzy Logic Theory,” *Color Research and Application*, vol. 33, no. 5, pp. 411–423, 2008.
- [15] H.R. Kang, *Computational Color Technology*, SPIE Press, 2006.
- [16] J.Z. Chang, J.P. Allebach, and C.A. Bouman, “Sequential Linear Interpolation of Multidimensional Functions,” *IEEE Transactions on Image Processing*, vol. 6, no. 9, pp. 1231–1245, 1997.
- [17] D. Kacker, A.U. Agar, J.P. Allebach, and B.J. Lucier, “Wavelet Decomposition based Representation of Nonlinear Color Transformations and Comparison with Sequential Linear Interpolation,” in *Proceedings of the IEEE International Conference on Image Processing*, 1998, pp. 186–190.
- [18] M.R. Gupta and R.M. Gray, “Color Conversions Using Maximum Entropy Estimation,” in *Proceedings of the IEEE International Conference on Image Processing*, 2001, pp. 118–121.
- [19] P. Hemingway, “n-Simplex Interpolation,” Tech. Rep. HPL-2002-320, HP, 2002, Available at <http://www.hp1.hp.com/techreports/2002/HPL-2002-320.html>.
- [20] H.R. Kang and P.G. Anderson, “Neural Network Applications to the Color Scanner and Printer Calibrations,” *Journal of Electronic Imaging*, vol. 1, no. 2, pp. 125–135, 1992.
- [21] S. Tominaga, “Color Notation Conversion by Neural Networks,” *Color Research and Application*, vol. 18, no. 4, pp. 253–259, 1993.
- [22] S. Usui, Y. Arai, and S. Nakauchi, “Neural Networks for Device-Independent Digital Color Imaging,” *Information Sciences*, vol. 123, no. 1/2, pp. 115–125, 2000.
- [23] M.J. Vrhel, “Approximation of Color Characterization MLUTS with Artificial Neural Networks,” in *Proceedings of the IEEE International Conference on Image Processing*, 2003, pp. 465–468.
- [24] E.W. Cheney, *Introduction to Approximation Theory*, AMS, Second edition, 2000.
- [25] J.-M. Muller, *Elementary Functions: Algorithms and Implementation*, Birkhäuser, Second edition, 2006.
- [26] W. Fraser, “A Survey of Methods of Computing Minimax and Near-Minimax Polynomial Approximations for Functions of a Single Independent Variable,” *Journal of the ACM*, vol. 12, no. 3, pp. 295–314, 1965.
- [27] K.N. Plataniotis and A.N. Venetsanopoulos, *Color Image Processing and Applications*, Springer-Verlag, 2000.
- [28] J.R. Kender, “Saturation, Hue, and Normalized Color: Calculation, Digitization Effects, and Use,” Tech. Rep., Carnegie Mellon University Computer Science Department, 1976.
- [29] S.E. Umbaugh, R.H. Moss, and W.V. Stoecker, “Automatic Color Segmentation of Images with Application to Detection of Variegated Coloring in Skin Tumors,” *IEEE Engineering in Medicine and Biology Magazine*, vol. 8, no. 4, pp. 43–52, 1989.
- [30] K.V. Mardia and P.E. Jupp, *Directional Statistics*, John Wiley & Sons, Second edition, 2001.